

REMARKS

Reconsideration of the present application is respectfully requested. Claim 17 has been canceled. Claims 1, 16 and 18 have been amended. Claims 19-23 are newly added. No new matter has been added.

Drawing Change

Figure 9 has been amended to correct a minor error. Specifically, the “yes” and “no” output branches of element 903 have been changed to “Reference only” and “Actual data block”, respectively. This change is fully supported by the original description. No new matter has been added.

Claim Rejections

Claims 1, 2 and 16 stand rejected under 35 U.S.C. § 102(b) based on U.S. Patent no. 6,728,898 of Tremblay et al. (“Tremblay”). Claims 3-15, 17 and 18 stand rejected under 35 U.S.C. § 103(a) based on Tremblay in view of U.S. Patent Application Publication no. 2002/0178355 of Selkirk et al. (“Selkirk”).

Section 102 Rejections

Before discussing the cited art and the claims, an overview and some background regarding the present invention may be helpful. The invention generally relates to a technique for mirroring data using two separate storage systems (e.g., storage servers) on a network, i.e., a source storage system and a destination storage system. The source storage system receives write requests from clients (for modifying stored data); however, the destination storage system does not receive the write requests from clients (see, e.g., Fig. 1 of the present application). The destination storage system determines how its mirror copy of data should be updated based only on information it receives from a source storage system. In other words,

the destination receives all modified data only from a source storage system. The modified data can be sent from the source storage system to the destination storage system in the form of log entries that reflect individual write transactions which the source storage system has received from clients (see, e.g., Fig. 6, blocks 605 – 607, of the present application).

Periodically, or according to some schedule, any data recently modified at the source storage system is saved to disk (or other persistent storage); this event is called a "consistency point" (CP) (see, e.g., present application at paragraph [0005]). Part of a CP event is updating the data at the destination storage system to reflect any recent changes, which is referred to as the "synchronization" phase of the CP event, or "CP sync".

In the known prior art, the CP sync phase involved sending over the network, from the source system to the destination system, all of the data modified since the last CP. In a system which handles large volumes of data, that would consume significant amounts of network bandwidth. See present application at paragraph [0006].

The present invention overcomes this problem. In particular, at least some of the modified data is not required to be sent from the source to the destination during the CP sync phase. Instead, only references to such data are sent to the destination during the CP sync phase. See, e.g., present application at paragraph [0020] et seq.

Accordingly, claim 1 recites:

1. (Currently amended) A method of mirroring data stored in a source storage system, the method comprising:
 - receiving at the source storage system a plurality of requests from a set of clients, the requests indicating modifications to be made to stored data;
 - saving modified data in the source storage system based on the requests;
 - receiving the modified data at a destination storage system from the source storage system, **wherein the destination storage system is configured to receive the modified data from the source storage system and not from any client of the set of clients**; and
 - during a synchronization phase, synchronizing data stored in the destination storage system with the data stored in the source storage system, including mirroring at least a portion of the modified data in the destination

storage system **without requiring said portion of the modified data to be sent from the source storage system to the destination storage system during the synchronization phase.** (Emphasis added.)

Tremblay fails to disclose or suggest such a method. In particular, Tremblay fails to disclose or suggest mirroring at least a portion of the modified data in the destination storage system without requiring said portion of the modified data to be sent from the source storage system to the destination storage system during the synchronization phase, where the destination storage system is configured to receive modified data from the source storage system and not from any client of the set of clients.

Tremblay discloses a method of mirroring data within a particular computer system. However, it is a completely different architecture from that of the present invention and, therefore, an entirely different approach to mirroring from the present invention. In Tremblay, two separate disks 105 and 110 are used to mirror data. Disks 105 and 110 are coupled to their I/O controllers 115 and 120, respectively. However, in Tremblay, the “client 125 . . . sends the same write requests 130 to both I/O controllers” (col. 3, lines 42-43)(emphasis added). Therefore, assuming *arguendo* Tremblay discloses a “source storage system” and a “destination storage system” as recited in claim 1, both the source and destination storage system receive the write requests from the client in Tremblay, in contrast with the present invention as recited in claim 1.

Because both the source and the destination storage system in Tremblay receive the write requests from the client, the data stored on disks 105 and 110 in Tremblay are normally synchronized (col. 3, line 48). The Tremblay system performs synchronization only to recover from some kind of fault or error (see col. 3, lines 63-67). In contrast, the present invention was developed for a system in which the destination storage system does not receive write requests from clients (as now recited in claim 1), and therefore, it must be frequently synchronized with the source system under normal operation. Because synchronization is not a normal event in

Tremblay, the Tremblay system does not experience the problem which the present invention was developed to solve, which is excessive consumption of network bandwidth due to the frequent transmission of modified data from source to destination during synchronization (see present application at paragraph [0006])(in fact, Tremblay does not even relate to a distributed system on a network).

Thus, Tremblay does not disclose or suggest the present invention as claimed, particularly, mirroring at least a portion of the modified data in the destination storage system without requiring said portion of the modified data to be sent from the source storage system to the destination storage system during the synchronization phase, where the destination storage system is configured to receive modified data only from the source storage system and not from any client of the set of clients. Moreover, because Tremblay is based on a different architecture and a completely different approach to mirroring than the present invention, the present invention also could not be obvious based on Tremblay.

Therefore, Applicants respectfully submit that the rejection of claim 1 is improper and should be withdrawn.

Claims 16 and 21 (new) include limitations similar to those discussed above and, therefore, are believed to be allowable for similar reasons to those discussed above.

Section 103 Rejections

Claims 3-15, 17 and 18 stand rejected under 35 U.S.C. § 103(a) based on Tremblay in view of Selkirk. The independent claims included in this rejection are claims 9, 13, 17 (canceled) and 18.

Claim 9 recites:

9. A method of mirroring data, the method comprising, in a first storage appliance:

receiving a plurality of requests to write data from a set of client devices, the requests for causing modification of a plurality of blocks of data stored in a first set of non-volatile storage devices coupled to the first storage appliance; storing modified data in the first set of non-volatile storage devices based on the requests;

initiating a process of synchronizing data in the first set of non-volatile storage devices with data stored in a second set of non-volatile storage devices coupled to a second storage appliance, including

sending each block of a first subset of the plurality of blocks from the first storage appliance to the second storage appliance, to cause the second storage appliance to store the blocks of the first subset in the second set of non-volatile storage devices, and

for each block of a second subset of the plurality of blocks, sending a reference from the first storage appliance to the second storage appliance, **instead of sending the corresponding block**, each said reference for use by the second storage appliance to locate the corresponding block in local storage of the second storage appliance and to store the corresponding block in the second set of non-volatile storage devices. (Emphasis added.)

The cited combination, Tremblay/Selkirk, fails to disclose or suggest all of the limitations of claim 9. Although Applicants arguments shall be directed to the alleged combination of references, it is useful to first consider their individual disclosures, in order to ascertain what combination, if any, could be made from them.

First, neither Tremblay nor Selkirk discloses or suggests sending each block of a first subset of the plurality of blocks from the first storage appliance to the second storage appliance, to cause the second storage appliance to store the blocks of the first subset in the second set of non-volatile storage devices. The Examiner cites Tremblay at col. 5, lines 20-20 as disclosing this; however, the cited section does not disclose or suggest that functionality. The system of Tremblay sends reference labels from the source to the destination, but it does not send the actual data blocks of the modified data. This is understandable, since there is no need in the Tremblay system to send any data from the source to the destination, because both the source and destination in Tremblay receive the data from the client. Likewise, Selkirk also fails to disclose or suggest this functionality of claim 9. For at least this reason, therefore, the rejection of claim 9 under section 103 was improper and should be withdrawn.

Second, claim 9 recites, “for each block of a second subset of the plurality of blocks, sending a reference from the first storage appliance to the second storage appliance, instead of sending the corresponding block . . .” The phrase “instead of” in here denotes that there is an alternative action that is performed (i.e., “the corresponding block is sent”) if a different condition is satisfied (as discussed in the preceding paragraph). But as noted above, in Tremblay, while reference labels are sent from source to destination during synchronization, the actual data blocks are never sent from source to destination system (at least so far as Tremblay discloses). Tremblay does not disclose or suggest any alternative to sending a reference during synchronization, and therefore, Tremblay fails to satisfy the “instead of” aspect of: “for each block . . . sending a reference from the first storage appliance to the second storage appliance, instead of sending the corresponding block . . .,” as recited in claim 9. Likewise, Selkirk also fails to disclose or suggest this functionality of claim 9. For this additional reason, therefore, the rejection of claim 9 under section 103 was improper and should be withdrawn.

Claim 18 includes limitations similar to those in claim 9 discussed above, and therefore, is patentable over the cited art for similar reasons. Note that the amendments to claim 18 are only to place that claim in what Applicants consider to be better form and not in response to the rejection, as Applicant does not believe any amendment was necessary for such purpose.

Dependent Claims

In view of the above remarks, a specific discussion of the dependent claims is considered to be unnecessary. Therefore, Applicants’ silence regarding any dependent claim is not to be interpreted as agreement with, or acquiescence to, the rejection of such claim or as waiving any argument regarding that claim.

Conclusion

For the foregoing reasons, the present application is believed to be in condition for allowance, and such action is earnestly requested.

If there are any additional charges/credits, please charge/credit our deposit account no. 02-2666.

Respectfully submitted,
BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP

Dated: May 17, 2006


Jordan M. Becker
Reg. No. 39,602

Customer No. 48102
12400 Wilshire Blvd.
Seventh Floor
Los Angeles, CA 90025
(408) 720-8300